

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 6

REMARKS

Claim rejections under 35 USC 112

Claim 1 has been rejected under 35 USC 112, second paragraph, as being indefinite. In particular, the Examiner states that it is not clear how the "one or more hidden output ports" can "receive an expanded port range." Although the Examiner has not said so, Applicant presumes that the Examiner is not sure about this limitation because he is not sure how *hidden* output ports could receive an expanded port range, since they are indeed *hidden*. In order to further along prosecution of this claim to allowance, Applicant has without prejudice amended claim 1 so that these hidden output ports are instead referred to as *other* output ports, and thus are not necessarily hidden. Therefore, it should be clear that *other* output ports could receive an expanded port range, since they are not necessarily hidden now.

Claim rejections under 35 USC 102

Claims 1-4, 8, and 10-20 have been rejected under 35 USC 102(e) as being anticipated by Huang (6,308,282). Applicant respectfully traverses this rejection. Claims 1, 11, and 16 are independent claims, from which claims 2-4, 8, 12-15, and 17-20 ultimately depend. Applicant submits that claims 1, 11, and 16 are not anticipated by Huang, such that the other claims are patentable for at least the same reasons.

Applicant specifically discusses claim 1 here as representative of independent claims 1, 11, and 16. That is, insofar as claim 1 is patentable over Huang, the other independent claims are patentable for at least the same reasons. This is now discussed in detail.

Why Huang as a general matter does not anticipate claim 1

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 7

Before getting into the specifics of claim 1, and why it is not anticipated by Huang, Applicant wishes to discuss the type of failover to which claim 1 is directed, and compare that to the type of failover to which Huang is directed. Applicant believes that this discussion will be insightful, because it is believed that there has been a fundamental misunderstanding of what the invention is directed to, as opposed to what Huang is directed to. That is, at a very fundamental level, Huang cannot anticipate the claimed invention.

The invention of claim 1 is directed to node failover. Thus, with respect to the first element of claim 1, where a first node of the network fails, the manager component selects an alternate address to route the destination address of the first node to a second node of the network. That is, the first node fails over to the second node. The first node no longer receives communications sent to the destination address, and instead the second node takes over for the first node, and receives communications sent to the destination address.

With respect to the second element of claim 1, where a third node of the network fails, the first switch remaps a destination address of a port for the third node to a port for a fourth node of the network. That is, the third node fails over to the fourth node. The third node no longer receives communications sent to the destination address, and instead the fourth node takes over for the third node, and receives communications sent to the destination address.

Finally, with respect to the third element of claim 1, when a fifth node of the network fails, the second switch remaps a destination address of an input port for the fifth node to an input port for the sixth node of the network. That is, the fifth node fails over to the sixth node. The fifth node no longer receives communications sent to the destination address, and instead the sixth node takes over for the fifth node, and receives communications sent to the destination address.

Therefore, it should be clear that claim 1 is directed to node failover. When a node fails, the invention in varying ways remaps the destination address of this node to another node. The node that has failed no longer can receive communications sent to the destination address, and is

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 8

basically out of the picture. The node to which the destination address was remapped takes over for the failed node, and thus receives communications sent to the destination address.

Huang is directed to a different type of failover, channel failover. A channel is described as follows in Huang.

The term channel is defined as the path from the network interface card of the sending node at one end to the network interface card of the receiving node at the other, and includes the drops, hubs, switches, bridges and bus. Two or more channels may share common sources, such as a[s] utilizing common hubs, switches, bridges, and bus. Alternatively, channel resources may be mutually exclusive. Two channels available on a node are referred to as dual channel, although a node is not limited to having only two channels.

(Col. 6, ll. 59-67) Thus, Huang is concerned with the failover of a path between two nodes to an alternate path between the two nodes. Neither of the nodes themselves actually fails, and more particularly, neither of the nodes is replaced by a failover node. That is, whereas the claimed invention is concerned with one node taking over for another node when the latter node fails, Huang is concerned with using a different path between two nodes when the current path between the two nodes fails. This is now discussed in detail in relation to Huang. Applicant tries to use the same portions of Huang that the Examiner relied upon in rejecting the claimed invention, so that it can be made clear that Huang talks about a different kind of failover than the claimed invention does.

First, Huang states the following:

[A]t least two nodes are each simultaneously connected to more than one network of a multiple-network system. Each node is provided with a network interface card for each connected network. Each node is further provided with an NIC switch capable of selecting one of the network interface cards. *A network fault-tolerance manager provides distributed detection of a failure on an active channel. The network fault-tolerance manager further provides failure recovery in manipulating the NIC switch of each sending node to select the network interface card to a stand-by channel. . . . In this embodiment, all nodes using the active channel are swapped to one stand-by channel upon detection of a failure on the active channel.* In a further embodiment, each node is connected to one active channel and one stand-by channel.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 9

(Col. 3, ll. 6-23) (Emphasis added) Thus, you can see what is going on here in Huang. Huang says that a node is connected to an active channel, or path, and a stand-by channel, or path. When the active channel of the node fails, the manager instructs the node to instead use the stand-by channel. The node itself does not fail – the channel (i.e., the network path) of the node fails. Another node does not take over for this node, as in the claimed invention, but rather the node just uses a different channel, or network path. Stated another way, while the claimed invention is limited to node failover in essence, Huang discloses channel or network path failover in essence.

Huang also goes into detail what is meant by a channel failing:

In reference to FIG. 1B and the preceding definitions, a failure of an active channel is defined as a failure of either network switch 240A, primary network bus 110A, any drop 130A or any network interface card 170A. Likewise, a failure of a stand-by channel is defined as a failure of either network switch 240B, secondary network bus 110B, any drop 130B or any network interface card 170B. Failures on an active channel will result in failure recovery, while failures on a stand-by channel will be reported without swapping devices or channels. *Note also that the definitions of active and stand-by are dynamic such that when an active channel fails and failure recovery initiates, the stand-by channel chosen for data traffic becomes an active channel.*

(Col. 8, ll. 21-33) (Emphasis added) Here, Huang describes what the failure of an active channel means, such that failover to a stand-by channel is performed. Basically, for an active channel between any two nodes 120 of FIG. 1B, any constituent component of device of the channel between the two nodes can fail, resulting in failover to a stand-by channel. Note that although a network interface card 170A of one of the nodes can fail as part of channel failure, the node itself does not fail. That is, another node does not take over for the node that had its card 170A fail. Rather, the node that had its card 170A fail has this channel failover to another channel that utilizes its other card 170B. Stated more clearly, whereas in the claimed invention one node takes over for another node when the latter node fails – such that the destination address of the former node is mapped to the latter node – in Huang, when a card 170A of one node fails, the channel (or network path) encompassing this card is failover to another channel using the other card 170B of

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 10

that same node. Another node does not ever take over for this node. The destination address of the node is not mapped to another node. The destination address remains mapped to the same node, and the node just receives communications on a different card 170B, which is part of another channel or network path.

Huang particularly describes two types of failures, "local" failures and "network" failures.

A local failure is described as follows.

A local failure in fault-tolerant network 200 is characterized by a device failure affecting communications to only one network interface card 170 of a node 120. For example, a local failure between node 120w and node 120z could be a failure of first switch 240₁ on the active channel connected to network interface card 170A of node 120w. *Swapping data communications to network interface card 170B of node 120w permits communication to node 120z through second switch 240₂.*

(Col. 8, ll. 52-63) (Emphasis added) So, Huang says you have two nodes, a first node 120w and a second node 120z. If there is a local failure in the channel between the first and the second nodes, such as a failure of the switch within the channel, you can simply failover to another channel between the first and the second nodes, where this other channel uses a different switch. Therefore, communication still is allowed to occur between the first and the second nodes 120w and 120z. By comparison, in the claimed invention, what happens is that, say, the first node fails in communication with the second node. Therefore, a third node takes over the destination address of the first node, and receives communications sent to this destination address by the second node. That is, there is no longer communication between the first and the second nodes, but instead there is no communication between the third and the second nodes, since the third node has taken over the destination address of the first node. Whereas Huang deals with channel failover, the claimed invention is limited to node failover.

The second type of failure in Huang, network failure, is described as follows.

A network failure in fault-tolerant network 200 is characterized by a device failure affect communications to all network interface cards of a node 120. For example, a network failure between node 120w and 120z could be a failure of second switch

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 11

240₂. . . . In this instance, network interface cards 170A and 170B of manager 250 close the ring, as shown by the dashed line 255, and allow data communications through manager node 250. With data communications served through manager node 250, *communications between node 120w and 120z are restored despite the failure of second switch 240₂.*

(Col. 8, l. 66, through col. 9, l. 13) (Emphasis added) So even in network failure, Huang still achieves channel failover, such that the communications between two nodes 120w and 120z can be restored. That is, Huang describes how communications between a first node 120w and a second node 120z are restored in light of a network failover, by channel failover. Compare this to the claimed invention: in the claimed invention, the first node, say, would fail. The destination address of the first node would be mapped to another node, which we will call the third node. Therefore, communications between the first node and the second node are not *restored* in the claimed invention, as in Huang, where the first node and the second node can continue to communicate with one another. Rather, the third node takes over for the first node (i.e., the first node fails over to the third node) in the claimed invention. After node failover, there are communications between the third and the second nodes in the claimed invention in this example, instead of between the first and the second nodes as in Huang. Again, this points to the basic difference between Huang and the claimed invention: Huang provides for the failover of a channel between two nodes when the channel fails, and the claimed invention provides for the failover of a node when the node fails.

At the end of the day, Huang provides for a very different kind of failover than the claimed invention does. In the claimed invention, one node takes over for another node by mapping the address of a failed node to a taking over node. In Huang, when a channel, or communication path, between two nodes fails, a stand-by channel, or stand-by communication path, takes over. There is no failure of a node in Huang, only failures of channels. This is indeed very nicely summed up in Huang itself:

In one embodiment, the network fault-tolerance manager switches communication of every network node from the channel experiencing failure to a stand-by channel. In another

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 12

embodiment, the network fault-tolerance manager switches communication of the node detecting a failure from the failed channel to a stand-by channel.

(Col. 18, ll. 40-46) There is just no node failover in Huang, and therefore as a general matter, Huang cannot anticipate the claimed invention.

Why the first element of claim 1 is specifically not anticipated by Huang

Applicant now discusses the first element of claim 1, which is “a manager component of a network having programmed therein alternate routes for a destination address, such that upon failure of a first node of the network to which the destination address is initially routed, the manager component selects one of the alternate routes to route the destination address to a second node of the network.” Let us first unpack this element a bit. A first node has a destination address routed to it. When this first node fails, the manager component instead routes the destination address to a second node. Thus, the second node takes over for the first node – or stated another way, the first node fails over to the second node. Communications sent to the destination address are thus received by the second node, instead of by the first node.

Now, the Examiner indicates that node 120I of FIG. 1B in Huang is the first node of the claimed invention, and node 120J of FIG. 1B in Huang is the second node. However, Huang is not concerned with the failure of the node 120I, such that the node 120J takes over for the node 120I. That is, Huang is not concerned with node failover, as in the claimed invention. Rather, Huang is concerned with the failure of a channel between the nodes 120I and 120J, such that the channel fails over to a standby channel, so that communications between the nodes 120I and 120J can still occur. For example, say the channel between the nodes 120I and 120J goes from the card 170A of the node 120I, to the switch 240A₁, and then to the card 170A of the node 120J. Now, say that the card 170A of the node 120I fails. What Huang tells you is that the channel that encompasses the card 170A of the node 120I should failover to the other channel connecting nodes 120I and 120J. That is, the standby channel between the nodes 120I and 120J should be

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 13

used, which goes from the card 170B of the node 120I, to the switch 240B₁, and then to the card 170B of the node 120J. What Huang does not do, as to which the claimed invention is limited, is having one node taking over for another node. Rather, Huang has one channel taking over for another channel. The two nodes still communicate with one another. A node is never replaced in Huang, only the channel to that node is replaced. That is, the manager component of Huang never routes the destination address of a first node to instead be routed to a second node, as in the claimed invention.

The manager component of Huang is specifically utilized in relation to network failures, as have been described above. That is, such network failures are described in Huang as follows:

A network failure in fault-tolerant network 200 is characterized by a device failure affect communications to all network interface cards of a node 120. For example, a network failure between node 120w and 120z could be a failure of second switch 240₂. . . . In this instance, network interface cards 170A and 170B of manager 250 close the ring, as shown by the dashed line 255, and allow data communications through manager node 250. With data communications served through manager node 250, *communications between node 120w and 120z are restored despite the failure of second switch 240₂.*

(Col. 8, l. 66, through col. 9, l. 13) (Emphasis added) Thus, in Huang, the manager component 250 allows communications between nodes 120w and 120z to remain even if the primary channel between them fails, such as the failure of the switch 240₂. The manager component does not replace the node 120w with the node 120z or with any other node. Huang does not achieve node failover; it achieves channel – or network path – failover.

Why the second element of claim 1 is specifically not anticipated by Huang

Applicant finally discusses the second element of claim 1, which is “a first switch of the network having a port for each of at least a third and a fourth node of the network, such that upon failure of the third node, the first switch remaps a destination address initially mapped to the port for the third node to the port for the fourth node.” Let us unpack this element a bit as well. A third node has a destination address routed to its port. When this third node fails, the first switch

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 14

remaps the destination address to a port of a fourth node. Thus, the fourth node takes over for the third node – or stated another way, the third node fails over to the fourth node. Communications sent to the destination address are thus received by the fourth node, instead of by the third node.

The Examiner indicates that the first switch is the switch 240A₁ of FIG. 1B in Huang, the third node is the node 120K, and the fourth node is the node 120X. However, as a first instance, this does not make sense. The switch 240A₁ is not directly connected to the node 120X at all. If you consider the cards 170A and 170B of the nodes 120 in FIG. 1B of Huang as the ports of the claimed invention, then there's no way the switch 240A₁ could remap a destination address from the port 170A of the node 120K to the port 170A or 170B of the node 120X.

Furthermore, the Examiner states that Huang anticipates this element of the claimed invention because “a manager node detects a failure of an active channel and updates/remaps the MAC address mapping table of the NIC switch to indicate that future communication should be directed to STANDBY node's address.” (Office Action, p. 3) However, this second element of the claimed invention has nothing to do with the manager component, which is what the first of the claimed invention, discussed above, is directed to. Rather, the second element of the claimed invention has the switch – not the manager component – mapping the port of one node for the port of another node.

Furthermore, there is no mapping of an address of one node to another node in Huang. One node does not take over for another node in Huang. Indeed, column 17, lines 28-44, as relied upon by the Examiner, states the following:

FIG. 15 depicts a flowchart of one embodiment . . . associated with a device swap failure recovery mode. If T_{skew} has expired for a sending node i at 1502, the receiving node checks at 1504 to see *if the failure is on the active channel or the stand-by channel*. If the failure is on the stand-by channel at 1504, the receiving node issues an alarm to the fault-tolerance manager at 1510 indicating failure of the stand-by channel. If the failure is on active channel at 1504, the MAC address mapping table of the NIC switch is updated at 1506. As shown in box 1520, *the update involves setting the active channel flag for node i*

Kashyap
Serial no. 09/917,314.
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 15

from A to B to indicate that future communication should be directed to node i's STANDBY address.

(Emphasis added) Thus, what Huang is saying here is that if the active channel between node i and a receiving node fails, then future communication should instead be directed to the standby channel (i.e., from the receiving node). That is, the receiving node is to direct future communication not to the active address of node i, but rather to the *standby* address of node i. The address of node i is not remapped to *another* node, as in the claimed invention. Rather, communications are directed to a *different address of the same node*. As has been described, Huang is not relevant to node failover, as in the claimed invention, but rather is relevant to active channel failover.

Claim rejections under 35 USC 103

Claims 5 and 6 have been rejected under 35 USC 103(a) as being unpatentable over Huang in view of Cao (6,618,371). Claim 7 has been rejected under 35 USC 103(a) as being unpatentable over Huang in view of Shah (6,694,361). Claim 9 has been rejected under 35 USC 103(a) as being unpatentable over Huang in view of Shah and further in view of Bolt (6,766,412). Claims 5-7 and 9 are dependent claims, depending ultimately from claim 1, and therefore are patentable for at least the same reasons that claim 1 is, as has been described above.


Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 16

Conclusion

Applicants have made a diligent effort to place the pending claims in condition for allowance, and request that they so be allowed. However, should there remain unresolved issues that require adverse action, it is respectfully requested that the Examiner telephone Applicants' Attorney so that such issues may be resolved as expeditiously as possible. For these reasons, this application is now considered to be in condition for allowance and such action is earnestly solicited.

Respectfully Submitted,



October 19, 2005
Date

Michael A. Dryja, Reg. No. 39,662
Attorney/Agent for Applicant(s)

Law Offices of Michael Dryja
704 228th Ave NE #694
Sammamish, WA 98074
tel: 425-427-5094, fax 206-374-2819